arXiv:1704.02906v1 [cs.CV] 10 Apr 2017

Multi-Agent Diverse Generative Adversarial Networks

Arnab Ghosh[1*] Viveka Kulharia[2*] Vinay Namboodiri[2] Philip H. S. Torr[1] Puneet K. Dokania[1]

[1] University of Oxford [2] IIT Kanpur

Abstract

This paper describes an intuitive generalization to the Generative Adversarial Networks (GANs) to generate samples while capturing diverse modes of the true data distribution. Firstly, we propose a very simple and intuitive multi-agent GAN architecture that incorporates multiple generators capable of generating samples from high probability modes. Secondly, in order to enforce different generators to generate samples from diverse modes, we propose two extensions to the standard GAN objective function. (1) We augment the generator specific GAN objective function with a diversity enforcing term that encourage different generators to generate diverse samples using a user-defined similarity based function. (2) We modify the discriminator objective function where along with finding the real and fake samples, the discriminator has to predict the generator which generated the given fake sample. Intuitively, in order to succeed in this task, the discriminator must learn to push different generators towards different identifiable modes. Our framework is generalizable in the sense that it can be easily combined with other existing variants of GANs to produce diverse samples. Experimentally we show that our framework is able to produce high quality diverse samples for the challenging tasks such as image/face generation and image-to-image translation. We also show that it is capable of learning a better feature representation in an unsupervised setting.

1. Introduction

Generating new data points given an unlabeled dataset is a core challenging problem of machine learning. In this paper we address this task using generative models. The underlying idea behind such models is to generate high-dimensional data such as images and texts using low-dimensional interpretable latent space. Though these models are highly useful in various applications, it is computationally challenging to efficiently train them as it normally requires intractable integration in a very high-dimensional complex space. Recently, there have been remarkable progress in this field with the development of generative models that do not explicitly require this integration and can be trained using back-propagation algorithm. Two famous examples of such models are Generative Adversarial Networks (GANs) [10] and Variational Autoencoders [13].

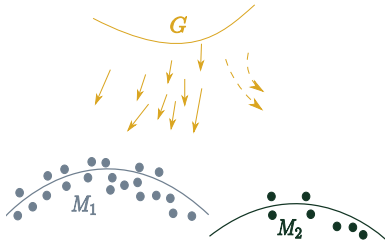

Figure 1. Intuitive visualization of the mode collapse problem. The gradients are biased towards the mode from which higher number of samples are drawn to form the real training data. Figure inspired by Che et al. [4].

In this work we focus on GANs as they are known to produce sharp and plausible images. Briefly, the objective function of GANs employ a generator and a discriminator where both are involved in a minimax game. The task of the discriminator is to learn the difference between 'real' (from true data distribution p_d) and 'fake' samples (from generator distribution p_g). However, the task of the generator is to maximize the mistakes of the discriminator. At convergence, the generator learns to produce real looking images. A few successful applications of GANs are video generation [25], image inpainting [20], image manipulation [27], 3D object generation [26], interactive image generation using few brush strokes [27], image super-resolution [14], diagrammatic abstract reasoning [8] and conditional GANs [18, 22].

Despite remarkable success of GANs in various applications, one of its major drawback is the problem of 'mode collapse' (as shown in the Figure 1) [2, 4, 5, 17, 23]. Even though, theoretically, at convergence the generator should be able to learn the true data distribution, practically this does not happen because of the minimax nature of the learn-

*Joint first author

ing problem which makes it very hard to reach the true equilibrium. Broadly speaking, there are two directions in which this issue can be resolved: (1) improving the learning aspect of GANs, similar to [2, 17, 23]; and (2) enforcing GANs to capture diverse modes, similar to [4, 5, 15]. In this work we focus on the latter.

Inspired by the multi-agent algorithm [1] and coupled GAN [15], we propose to use multiple generators with one discriminator, and allow generators to share information. We call this the Multi-Agent GAN architecture as shown in the Figure 2. In more detail, similar to the standard GAN, the objective of each generator here is to maximize the mistakes of the common discriminator. To allow different generators to share information with each other, we share the major chunk of parameters among all the generators. Another reason behind sharing these parameters is the fact that initial layers capture high-frequency structures which is almost the same for a particular type of dataset (for example, faces), therefore, sharing them reduces redundant computations by different generators. Now, to make different generations visually different, we propose to use different end-layers for each generator. Naively using this simple approach may lead to the *trivial solution* where all the generators learn to generate multiple but *similar* samples. To resolve this issue and to generate different visually plausible samples capturing diverse high probability modes of the true data distribution, we propose two extensions to the standard GAN. First extension involves augmenting the generator specific objective function with a diversity enforcing term where the diversity comes from a user-defined task-specific function. For example, the diversity enforcing term can be a similarity measure between different samples generated by different generators. This term ensures that different generations are dissimilar while capturing high-probability regions. The second extension involves modifying the discriminator objective function. In this case, along with finding the real and the fake samples, the discriminator has to correctly predict the generator that generated the given fake sample. Intuitively, in order to succeed in this task, the discriminator must learn to push generations corresponding to different generators towards different identifiable modes. Combining the Multi-Agent GAN architecture with the diversity enforcing terms allows us to generate diverse plausible samples, thus the name Multi-Agent Diverse GAN (MAD-GAN).

We show the effectiveness of our framework on both synthetic and real-world datasets. We use a synthetic dataset to show that our framework is able to capture diverse clustered modes even if the training data corresponding to some modes are not enough compared to the other modes, a situation highly likely in high-dimensional real world data. Next, we show diverse generations on challenging tasks of *image to image translation* [11] and *image/face generation* [5, 21]. In the end, using the SVHN dataset [19], we show that our framework is capable of learning a better feature representation in an unsupervised setting.

2. Related Work

The recent work called InfoGAN [5] proposed an information-theoretic extension to GANs in order to address the problem of mode collapse. Briefly, InfoGAN disentangles the latent representation by assuming a factored representation of the latent variables. In order to enforce the generator to learn factor specific generations, InfoGAN maximizes the mutual information between the factored latents and the generator distribution. Che et al. [4] proposed a mode regularized GAN (ModeGAN) which uses an encoder-decoder paradigm. The basic idea behind ModeGAN is that if a sample from the true data distribution p_d belongs to a particular mode, then the sample generated by the generator (fake sample) when the true sample is passed through the encoder-decoder is likely to belong to the same mode. ModeGAN assumes that there exists enough true samples from a mode for the generator to be able to capture it. Another work by Metz et al. [17] proposed a surrogate objective for the update of the generator with respect to the unrolled optimization of the discriminator (UnrolledGAN) to address the issue of convergence of the training process of GANs. This improves the training process of the generator which in turn allow the generators to explore wide coverage of the true data distribution.

Liu et al. [15] presented Coupled GAN, a method for the training of two generators with shared parameters to learn joint distribution of the data. The shared parameters guide both the generators towards similar subspaces but since they are trained independently on two domains, they promote diverse generations. Durugkar et al. [7] proposed a model with multiple discriminators whereby an ensemble of multiple discriminators have been shown to stabilize the training of the generator by guiding it to produce better samples. In terms of employing multiple generators, our work is closest to [15, 9]. However, while using multiple generators, our method explicitly enforces them to capture diverse modes.

3. Preliminaries

As our objective is to efficiently capture diverse modes for image generation using Generative Adversarial Networks (GANs) [10], we start with a brief review of GANs. Given a set of unlabelled samples $\mathcal{D} = (x_i)_{i=1}^n$ from the true data distribution p_d, the GAN learning problem finds the optimal parameters θ_g of a generator $G(z; \theta_g)$ that allows us to sample from an approximate data distribution p_g. Here $z \sim p_z$ is the prior input noise sampled from the distribution p_z (for example, a normal distribution). In order to learn the optimal parameters θ_g, the GAN objective

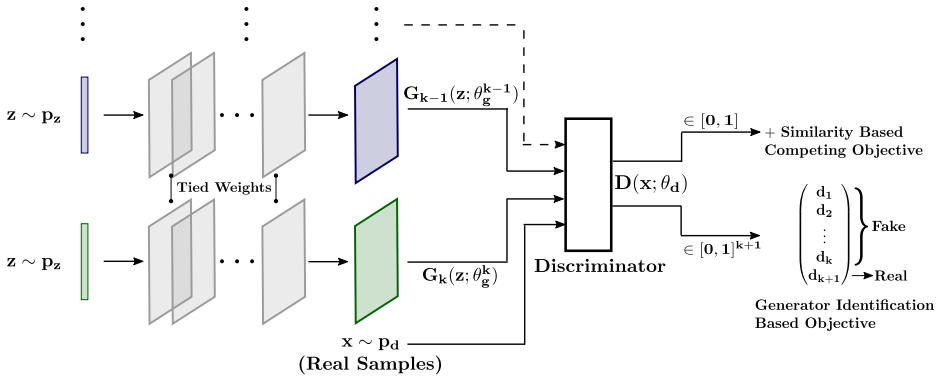

Figure 2. The Multi-Agent Diverse GAN (MAD-GAN) architecture. All the generators share parameters of all the layers except the last one. Two proposed diversity enforcing objectives, 'competing' and 'generator identification', are shown at the end of the discriminator.

function (Equation 1) employs a discriminator $D(x;\theta_d)$ that learns whether the sample x is 'real' (sampled from p_d) or 'fake' (sampled from p_g). The overall GAN objective function is as follows:

$$\min_{\theta_g}\max_{\theta_d} V(\theta_d,\theta_g) := \mathbb{E}_{x\sim p_d}\log D(x;\theta_d) + \mathbb{E}_{z\sim p_z}\log\big(1-D(G(z;\theta_g);\theta_d)\big) \quad (1)$$

The above objective function is optimized in a block-wise manner where θ_d and θ_g are optimized one at a time while fixing the other. Intuitively, for a given sample x (either sampled from p_d or p_g) and the parameter θ_d, the function $D(x;\theta_d) \in [0,1]$ produces a score that represents the probability of x belonging to the true data distribution p_d (or probability of it being real). Hence, the discriminator in the above objective function is trying to learn parameters θ_d that maximizes this score for the true samples (from p_d) while minimizing it for the fake samples $\tilde{x} = D(z;\theta_g)$ (from p_g). On the other hand, the generator is minimizing the function $\mathbb{E}_{z\sim p_z}\log\big(1-D(G(z;\theta_g);\theta_d)\big)$, equivalently maximizing $\mathbb{E}_{z\sim p_z}\log D(G(z;\theta_g);\theta_d)$. Thus, the generator is learning to maximize the scores for the fake samples (from p_g), which is exactly opposite to what discriminator is trying to achieve. In this manner, the generator and the discriminator are involved in a minimax game where the task of the generator is to maximize the wrong predictions of the discriminator. At convergence, the generator learns to generate real samples, which means $p_g = p_d$.

4. Multi-Agent Diverse GAN

The task of the generator in the GAN objective function is much harder than the task of the discriminator as it has to produce real looking images while maximizing the mistakes of the discriminator. This, along with the minimax nature of the objective function raise several challenges with GANs such as [2, 4, 5, 17, 23]: (1) mode collapse; (2) difficult optimization; and (3) trivial solution. In this work we propose a new framework to address the first challenge of 'mode collapse' by increasing the capacity of the generator while we use the well known optimization tricks to partially avoid other challenges [2].

Briefly, we propose a *Multi-Agent GAN architecture* that employs multiple generators in order to generate different samples while capturing high probability regions of the true data distribution. In order to encourage different generators to move towards different diverse modes we propose two extensions to the standard GAN – (1) *similarity based competing objective*: augment the GAN objective function with a diversity enforcing term that ensures that the generations from different generators are diverse where the diversity depends on a user-defined task-specific function and; (2) *generator identification based objective*: modify the objective function of the discriminator where along with finding the fake samples, the discriminator has to find the generator that produced the given fake sample. Combining the multi-agent architecture along with the diversity enforcing terms leads to a framework capable of capturing diverse high probability regions of the true data distribution. Our proposed framework is very general in the sense that it can be used with different variants of GANs. In what follows we describe our framework in detail.

4.1. Multi-Agent GAN Architecture

In this section we describe our proposed architecture as shown in Figure 2. It involves k generators and one discriminator. We allow all the generators to share information by tying the parameters of the initial layers. A naive way would be to allow different generators to have entirely different set of parameters. However, since initial layers of a generator capture high-frequency structures which is almost the same for a particular type of dataset (for example, faces), therefore, parameter sharing is essential to avoid redundant computations and allow different generators to converge faster. In order to produce different visually looking generations from different generators, we use completely independent

set of parameters for the end-layer of each generator.

Specifically, given $z \sim p_z$ for the i-th generator, similar to the standard GAN, the first step involves generating a sample (for example, an image) $\tilde{x}_i$. Notice that each generator receives a different latent input sampled from the same distribution. Naively using this simple approach may lead to the trivial solution where all the generators learn to generate multiple *similar* samples. In what follows we discuss two methods to avoid this issue by enforcing diversity among the generated samples.

4.2. Enforcing Diverse Modes

Here we propose two alternatives in order to push different generators towards different diverse modes. Figure 3 presents an intuitive way to show the effect of enforcing diversity over generators.

4.2.1 Similarity based competing objective

The approach presented here is motivated by the fact that the samples from different modes must look different. For example, in the case of images, these samples should differ in terms of texture, color, shading, and various other cues. Thus, different generators must generate dissimilar samples where the dissimilarity comes from a task-specific function. Before delving into the details, let us first define some notations in order to avoid clutter. We denote θ_g^i as the parameters of the i-th generator. The set of generators is denoted as $K = \{1, \cdots, k\}$. Given random noise z to the i-th generator, the corresponding generated sample $G_i(z; \theta_g^i)$ is denoted as $g_i(z)$. Using these notations and following the above discussed intuitions, we impose following constraints over the i-th generator while updating its parameters:

$$D(G_i(z;\theta_g^i);\theta_d) \geq D(G_j(z;\theta_g^j);\theta_d) + \Delta\big(\phi(g_i(z)), \phi(g_j(z))\big), \quad \forall j \in K \setminus i \tag{2}$$

where, $\phi(g_i(z))$ denotes the mapping of the generated image $g_i(z)$ by the i-th generator into a feature space and $\Delta(.,.) \in [0, 1]$ is the similarity function. Higher the value of $\Delta(.,.)$ more similar the arguments are. Intuitively, the above set of constraints ensures that the discriminator score for each generator should be higher than all other generators with a margin proportional to the similarity score. If the samples are similar, the margin increases and the constraints become more active. We use unsupervised learning based representation as our mapping function $\phi(.)$. Precisely, given a generated sample $g_i(z)$, $\phi(g_i(z))$ is the feature vector obtained using the discriminator of our framework. This is motivated by the feature matching based approach to improve the stability of the training of GANs [23]. The $\Delta(.,.)$ function used in this work is the standard cosine similarity based function. The above mentioned constraints can be satisfied by maximizing an equivalent unconstrained objective function as defined below:

$$U(\theta_g^i, \theta_d) := f\Big(D(G_i(z;\theta_g^i);\theta_d) - \frac{1}{k-1}\sum_{j \in K \setminus i}\big(D(G_j(z;\theta_g^j);\theta_d) + \Delta(\psi_i, \psi_j)\big)\Big) \tag{3}$$

where, $f(a) = \min(0, a)$, $\psi_i = \phi(g_i(z))$, and $\psi_j = \phi(g_j(z))$. Intuitively, if the argument of $f(.)$ is positive, then the desirable constraint is satisfied and there is no need to do anything. Otherwise, maximize the argument with respect to θ_g^i. Note that instead of using all the constraints independently, we use the average of all of them. Another approach would be to use the constraint corresponding to the j-th generator that maximally violates the set of constraints shown in Equation 2. Experimentally we found that the training process of the average constraint based objective is more stable than the maximum violated constraint based objective. The intuition behind using these constraints comes from the well know 1-slack formulation of the structured SVM framework [12, 24]. Thus, the overall objective for the i-th generator is:

$$\min_{\theta_g^i} \; V(\theta_d, \theta_g^i) - \lambda\, U(\theta_g^i, \theta_d) \tag{4}$$

where $\lambda \geq 0$ is the hyperparameter. Algorithm 1 shows how to compute gradients corresponding to different generators for the above mentioned objective function. Notice that, once sampled, the same z is passed through all the generators in order to enforce constraints over a *particular generator* (as shown in Equation 2). However, in order for constraints to not to contradict with each other while updating another generator, a different z is sampled again from the p_z. The Algorithm 1 is shown for the batch of size one which can be trivially generalized for any given batch sizes. In the case of discriminator, the gradients will have exactly the same form as the standard GAN objective. The only difference is that in this case the fake samples are being generated by k generators, instead of one.

4.2.2 Generator identification based objective

Inspired by the discriminator formulation for the semi-supervised learning [23], we use a generator identification based objective function that, along with minimizing the score $D(\tilde{x}; \theta_d)$, requires the discriminator to identify the generator that generated the given fake sample $\tilde{x}$. In order to do so, as opposed to the standard GAN objective function where the discriminator outputs a scalar value, we modify it to output $k+1$ soft-max scores. In more detail, given the set of k generators, the discriminator produces a soft-max probability distribution over $k+1$ classes. The score at $(k+1)$-th index represents the probability that the sample belongs to

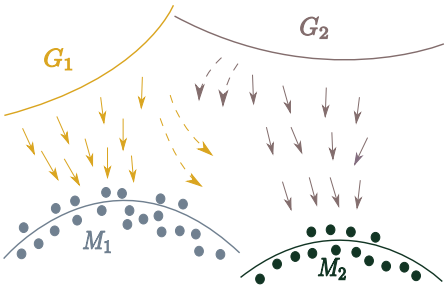

Figure 3. Intuitive visualization of how different generators are pushed towards different high probability modes as explained in Section 4.2. Note that M_1 and M_2 could be a cluster of modes where each cluster themselves contain different modes. The arrows abstractly represent generator specific gradients for the purpose of building intuition.

the true data distribution and the score at $j \in \{1, \ldots, k\}$-th index represents the probability of it being generated by the j-th generator. Under this setting, while learning θ_d, we optimize the cross-entropy between the soft-max output of the discriminator and the Dirac delta $\delta \in \{0, 1\}^{k+1}$, where for $j \in \{1, \ldots, k\}$, $\delta(j) = 1$ if the sample belongs to the j-th generator, otherwise $\delta(k+1) = 1$. Therefore, the objective function (1) for the optimization of θ_d while keeping θ_g constant becomes:

$$\max_{\theta_d} \mathbb{E}_{x \sim p} H(\delta, D(x; \theta_d)) \tag{5}$$

where, $Supp(p) = \cup_{i=1}^{k} Supp(p_{g_i}) \cup Supp(p_d)$ and $H(.,.)$ is the negative of the cross entropy function. Intuitively, in order to correctly predict which generator produced a given fake sample, the discriminator must learn to push different generators towards different identifiable modes. To update parameters, the gradient for each generator is simply $\nabla_{\theta_g^i} \log(1 - D(G_i(z; \theta_g^i); \theta_d)))$. Notice that all the generators in this case can be updated in parallel. For the discriminator, given $x \sim p$ (can be real or fake) and corresponding δ, the gradient is $\nabla_{\theta_d} \log D_j(x; \theta_d)$, where $D_j(x; \theta_d)$ is the j-th index of $D(x; \theta_d)$ for which $\delta(j) = 1$. Therefore, using the generator identification based objective function requires very minor modifications on the standard GAN optimization algorithm and can be easily used for different variants of GAN. Using Theorem 1 we show that the optimal generator learns the true data distribution in the form of a mixture of distributions $\frac{1}{k} \sum_{i=1}^{k} p_{g_i}$ where every distribution component is assigned an equal weight of $\frac{1}{k}$.

Proposition 1. *For fixed generators, the optimal distribution learned by the discriminator D has following form:*

$$D_{k+1}(x) = \frac{p_d(x)}{p_d(x) + \sum_{i=1}^{k} p_{g_i}(x)}, \tag{6}$$

$$D_i(x) = \frac{p_{g_i}(x)}{p_d(x) + \sum_{i=1}^{k} p_{g_i}(x)}, \forall i \in \{1, \cdots, k\}.$$

where, $D_i(x)$ represents the i-th index of $D(x; \theta_d)$, p_d the true data distribution, and p_{g_i} the distribution learned by the i-th generator.

Proof. For fixed generators, the objective function of the discriminator is to maximize

$$\mathbb{E}_{x \sim p_d} \log D_{k+1}(x) + \sum_{i=1}^{k} \mathbb{E}_{x_i \sim p_{g_i}} \log D_i(x_i) \tag{7}$$

where, $\sum_{i=1}^{k+1} D_i(x) = 1$ and $D_i(x) \in [0, 1], \forall i$. The above equation can be written as:

$$\int_x p_d(x) \log D_{k+1}(x) dx + \sum_{i=1}^{k} \int_x p_{g_i}(x) \log D_i(x) dx$$

Using $\sum_{i=1}^{k+1} D_i(x) = 1$ and equating the derivative w.r.t. $D_i(x)$ to be zero we obtain the following form

$$D_i(x) = \frac{p_{g_i}(x)\Big(1 - \sum_{j=1, j \neq l}^{k+1} D_j(x)\Big)}{p_{g_l}(x)} \tag{8}$$

Note that in the above equation we abuse notation by assuming $p_{g_{k+1}} := p_d$ to avoid clutter. Equation (8) holds for any $l \in [1, \ldots, k+1]$ and $i \neq l$. Using $\sum_{i=1}^{k+1} D_i(x) = 1$ and $l = k+1$ we obtain

$$D_{k+1}(x) = \frac{p_d(x)}{p_d(x) + \sum_{i=1}^{k} p_{g_i}(x)} \tag{9}$$

Putting $D_{k+1}(x)$ into Equation (8) with $l = k+1$ we obtain

$$D_i(x) = \frac{p_{g_i}(x)}{p_d(x) + \sum_{i=1}^{k} p_{g_i}(x)}, \forall i \in \{1, \cdots, k\}. \tag{10}$$

□

Theorem 1. *Given the optimal discriminator, the objective function for the training of the generators boils down to minimizing*

$$KL\Big(p_d(x) || p_{avg}(x)\Big) + kKL\Big(\frac{1}{k} \sum_{i=1}^{k} p_{g_i}(x) || p_{avg}(x)\Big)$$
$$- (k+1) \log(k+1) + k \log k \tag{11}$$

where, $p_{avg}(x) = \frac{p_d(x) + \sum_{i=1}^{k} p_{g_i}(x)}{k+1}$. The above objective function obtains its global minimum if $p_d = \frac{1}{k} \sum_{i=1}^{k} p_{g_i}$ with the objective value of $-(k+1) \log(k+1) + k \log k$.

Algorithm 1 Updating Generators for Competing Objective

input θ_d; $p(z)$; $\theta_g^i, \forall i \in \{1, \cdots, k\}$; λ.

for each generator $i \in \{1, \cdots, k\}$ **do**
  Sample noise from the given noise prior $z \sim p_z$.
  Obtain the generated sample $G_i(z; \theta_g^i)$ and corresponding feature vector $\psi_i = \phi(G_i(z; \theta_g^i))$.
  $\nu \leftarrow 0$.
  for each generator $j \in \{1, \cdots, k\} \setminus i$ **do**
    Compute feature vector $\psi_j = \phi(G_j(z; \theta_g^j))$.
    $\nu \leftarrow \nu + D(G_j(z; \theta_g^j); \theta_d) + \Delta(\psi_i, \psi_j)$.
  end for
  $\nu \leftarrow D(G_i(z; \theta_g^i); \theta_d) - \frac{\nu}{k-1}$.
  if $\nu \geq 0$ **then**
    $\nabla_{\theta_g^i} \log(1 - D(G_i(z; \theta_g^i); \theta_d)))$.
  else
    $\nabla_{\theta_g^i} \left(\log(1 - D(G_i(z; \theta_g^i); \theta_d))) - \lambda U(\theta_g^i, \theta_d) \right)$.
  end if
end for

output

Proof. The objective function of the set of generators is to minimize the following:

$$\mathbb{E}_{x \sim p_d} \log D_{k+1}(x) + \sum_{i=1}^{k} \mathbb{E}_{x \sim p_{g_i}} \log(1 - D_{k+1}(x)) \tag{12}$$

Using Proposition 1 we obtain

$$\begin{aligned} &\mathbb{E}_{x \sim p_d} \log \left[\frac{p_d(x)}{p_d(x) + \sum_{i=1}^{k} p_{g_i}(x)} \right] + \\ &\sum_{i=1}^{k} \mathbb{E}_{x \sim p_{g_i}} \log \left[\frac{\sum_{i=1}^{k} p_{g_i}(x)}{p_d(x) + \sum_{i=1}^{k} p_{g_i}(x)} \right] \\ &:= \mathbb{E}_{x \sim p_d} \log \left[\frac{p_d(x)}{p_{avg}(x)} \right] + k \mathbb{E}_{x \sim p_g} \log \left[\frac{p_g(x)}{p_{avg}(x)} \right] \\ &- (k+1) \log(k+1) + k \log k \end{aligned} \tag{13}$$

where, $p_g = \frac{\sum_{i=1}^{k} p_{g_i}}{k}$, $Supp(p_g) = \bigcup_{i=1}^{k} Supp(p_{g_i})$, and $p_{avg}(x) = \frac{p_d(x) + \sum_{i=1}^{k} p_{g_i}(x)}{k+1}$. Note that Equation (13) is exactly the same as Equation (11) which obtains its minimum value when $p_d = \frac{\sum_{i=1}^{k} p_{g_i}}{k}$. In the case of one generator ($k = 1$), Equation (11) obtains the Jensen-Shannon divergence based objective function as shown in [10] with the minimum objective value of $-\log 4$. □

5. Experiments

We show the efficacy of our framework on synthetic and real-world datasets. First of all, using a toy experiment we show that our framework is able to capture diverse clustered modes even if the training data corresponding to a particular cluster is not enough compared to other clusters. This situation is highly likely in high-dimensional real world data (images) as they lie in a very complex space with large number of clustered modes. Next, we show the effectiveness of our framework on challenging tasks of *image to image translation* [11], *image/face generation*, and *unsupervised representation learning*. In more detail, corresponding to the real-world datasets, the first task that we consider is the image to image translation where given an input image, the objective is to map it to another image. For example, given an image of a scene taken in the night, the objective is to generate corresponding day images. The second task we consider is the face/image generation using CelebA [16] and ImageNet [6] datasets. For the face generation, we show results using both InfoGAN [5] and DC-GAN [21]. We also show image generations using the very challenging ImageNet dataset for which we again employ DC-GAN. In the end, we show how the diversity promoting objectives and the multiple generators can help in improving the representation learning of the discriminator in unsupervised setting which can then be used for downstream tasks such as semi supervised digit identification. With these experiments we show that the generations using our framework are highly plausible and diverse.

Note that in all our real-data experiments, CNN architectures, datasets, and the parameters used are exactly the same as provided in their respective references. For the sake of

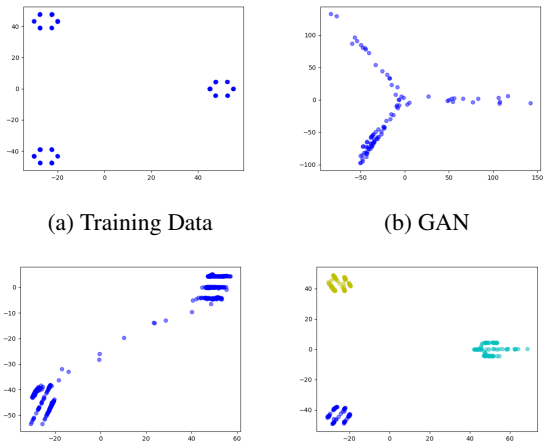

(a) Training Data (b) GAN (c) GAN+BN (d) Our (MAD-GAN)

Figure 4. A toy example to show how our method can capture diverse clustered modes. Figure 4a shows the real training data. Each dot itself represents many data points. Figures 4b and 4c shows the modes captured by GAN (600 epochs) and GAN with batch normalization (600 epochs), respectively. Finally, Figure 4d shows the modes captured by our method (200 epochs) where different colors represent different generators.

completeness, we provide these details, including the ones for the toy example, in the supplementary material. We use three generators in all our experiments and show results using 'competing' and 'generator identification' based objective functions. We experimented with different values of λ used only in 'competing objective' (Equation 4) and found 10^{-3} and 10^{-4} to perform very well for the image to image translation task and DC-GAN related tasks, respectively.

5.1. Capturing clustered Modes: Toy Example

We generate synthetic data with three clustered modes. To do so, we use GMM [3] with 18 Gaussians and select their means such that it forms three clusters, each consisting of 6 Gaussians (Figure 4a). We ensure that for one particular cluster we generate less amount of data in order to imitate the data imbalance situation. This situation is highly probable in high-dimensional real world datasets. Under this setting, Figure 4b shows the modes captured by the standard GAN. Clearly, GAN has collapsed all the modes, which is expected. Figure 4c shows the modes captured by GAN with batch normalization. Batch normalization allows better training of the parameters which in turn allows GANs to capture multiple modes. However, it is shown in the toy experiment that, though GAN with batch normalization is able to focus on two clusters (fails to capture the cluster with less data), the distribution is distorted and does not match the train data distribution. Finally, Figure 4d shows the modes captured by our method using 'generator identification' based diversity term. Different colors are used to represent modes corresponding to different generators. It is evident from this experiment that our method is able to capture all the clustered modes which includes modes for which the training data was very limited. Also, the generated distribution is similar to the train data distribution.

5.2. Image to Image Translation

In this section we show results on the challenging task of image to image translation [11] which uses patch based Conditional GANs [18]. Similarly, we also employ patch based conditional GAN for solving this task. We follow the same approach as presented in [11]. We show diverse and highly appealing results using both the diversity promoting objectives: 'competing' and 'generator identification'. Recall that in the case of competing objective we use cosine based similarity to enforce diverse generations, an important criteria for real images. We show results for the following two situations where diverse solution is useful: (1) given sketches of bags, generate real looking bags; and (2) given night images of places, generate their equivalent day images. Figures 5 and 6 show generated results for them. We clearly notice that each generator is able to produce meaningful and diverse images. The generated bags, as shown in Figure 5, differ in terms of texture, color, and design patterns. Similarly, the generated day images in Figure 6 differ in terms of lighting conditions, sky patterns, weather conditions, and many other minute yet useful cues.

5.3. Face generation with InfoGAN as the generator

Given the CelebA face dataset, the objective is to generate diverse faces. Here we replace each generator in our framework with InfoGAN [5] and analyze the effect of our diversity enforcing objective ('generator identification'). The generations are shown in the Figure 7. Note that the InfoGAN itself can disentangle various factors of variations upto certain extent using categorical variables. We refer reader to [5] for further details about InfoGAN. Thus, combining different InfoGANs with our 'generator identification' based objective pushes each of them towards different clustered modes (as described in the toy experiment in Section 5.1). Since each clustered mode is now captured by individual InfoGAN, varying the categorical variables of a given InfoGAN allows us to navigate over different modes of a clustered mode. Another interesting fact emerged that none of the factors of variation were shared between different InfoGANs, thus, proving an elegant way of pattern mining. It has to be noted though that the 'competing objective' in this case did not provide as appealing results as provided by the 'generator identification'. We suspect that this is because of the fact that it depends on the feature vectors of the generated samples to obtain the similarity measure. In this case, the feature space we used may not be discriminative enough to small facial variations, thus, did not provide plausible results. We believe that a more specific similarity function would provide better results and we intend to explore this in the future.

5.4. Face/Image generations with DC-GAN

In this section we show face (CelebA dataset) and image generations (ImageNet dataset) using DC-GAN [21] as our generators. Again, we use the same setting as provided in DC-GAN. The results for the face and the image generations are shown in Figures 8 and 9, respectively. As depicted in 8a, with the generator identification based objective, the first generator is generating faces as the normal DC-GAN architecture would. The second generator is generating faces like oil-paintings, while the third generator seems like it is generating images with coarse strokes without much focus on the background. Similarly, in Figure 8b, with the similarity based competing objective, the first generator generates as expected by the DC-GAN, while the second generator generates with high detailing of the faces and exaggerated facial features, and the third generator generates images with sketch like quality with an artistic touch. The images generated when trained over ImageNet dataset, as shown in Figure 9, are able to capture some meaningful patterns corresponding to birds, insects, sky, water, among

(a) Competing Objective (b) Generator Identification

Figure 5. Diverse generations for 'sketch-bag to real-bag' generation task. The first column in each subfigure represents the input. The remaining three columns show the diverse outputs of different generators. It is evident from that different generators are able to produce very diverse results capturing different colors (blue, green, black etc), textures (leather, fabric, plastic etc), design patterns, among others.

(a) Competing Objective (b) Generator Identification

Figure 6. Diverse generations for 'night to day' image generation task. First column in each subfigure represents the input. The remaining three columns show the diverse outputs of different generators. It is evident from that different generators are able to produce very diverse results capturing different lighting conditions, different sky patterns (cloudy vs clear sky), different weather conditions (winter, summer, rain), different landscapes, among many other minute yet useful cues.

others. However, though these results are better than the existing approaches, we must accept that they are not very realistic because of the fact that training over ImageNet dataset, which contains huge variations in terms of objects, texture, weather, etc., is highly challenging. To the best of our knowledge, no existing generative model is able to produce very promising results on this dataset.

5.5. Unsupervised Representation Learning

Similarly to DC-GAN [21], we train our framework using SVHN dataset [19]. The trained discriminator is then used to extract features. Using these features, we train an SVM for the classification task. For the 'competing' and 'generator identification' based objectives we obtain misclassification error of 18.3% and 17.5%, respectively. This

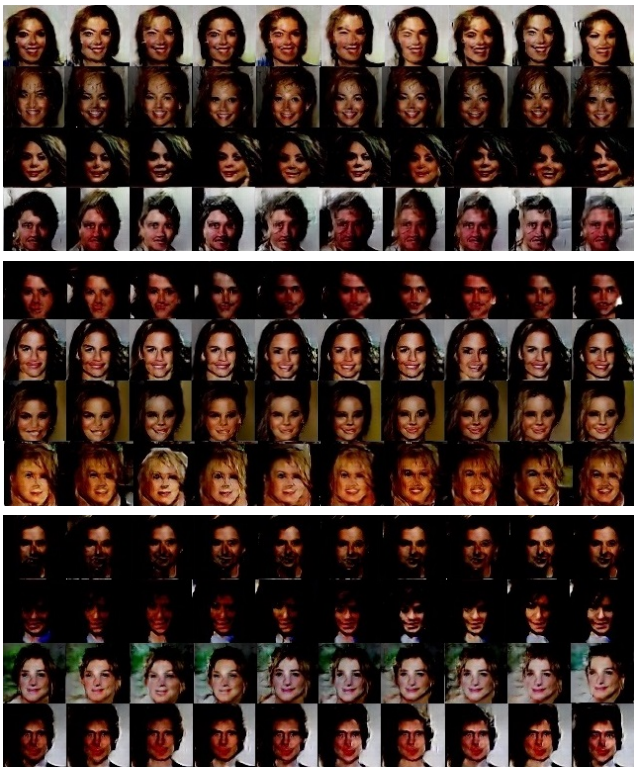

Figure 7. **Face generations using InfoGANs as generators:** Each row represent generations corresponding to a random input and columns represent variations in the generations for different categorical codes of the InfoGAN. We used three generators for this experiment. It is clear that the first generator (first four rows) capture variations in the hair in terms of style, quantity, and color. The second generator (rows 5 to 8) captures variations in mustaches (on male face) and emotions (such as smile). Similarly, the third generator (rows 9 to 12) captures varying forms of lighting on a face from the top/below/right of the face, and varying colors of a particular male face.

is almost 5% better than the results reported by DC-GAN (22.48%), which clearly indicates that our framework is able to learn better feature space in an unsupervised setting, which is a highly challenging and useful task.

6. Conclusion

We presented a framework, Multi-Agent Diverse GAN, for generating diverse and meaningful samples using multiple generators and diversity enforcing terms. Through various challenging experiments such as image to image translation and image/face generation, we showed that our framework is capable of capturing diverse and meaningful modes. Using toy experiment we showed that even in situations where there exist imbalance in the training data, our framework is able to capture modes for which we had very limited data.

(a) Generator Identification

(b) Competing Objective

Figure 8. **Face generations using DC-GANs as generators:** Each row represents a generator. Each column represents generations for a given random noise input z. Figure 8a shows results corresponding to the generator identification based objective and Figure 8b corresponding to the similarity based competing objective. Clearly, the generations are highly plausible and diverse compared to the standard DC-GAN results.

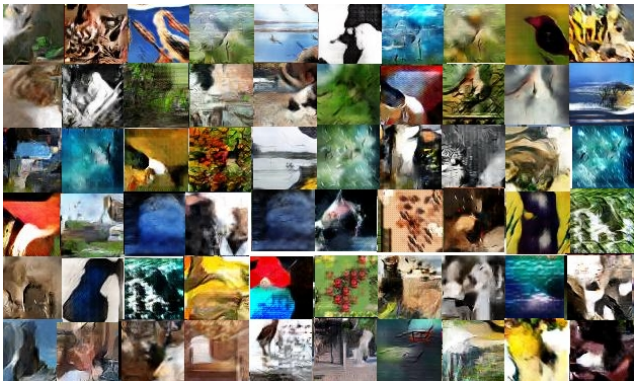

Figure 9. **Image generations using DC-GANs as generators:** Image generations using our framework when trained over the challenging ImageNet dataset. Each generator in this case is a DC-GAN. The results are shown using 'generator identification' based diversity enforcing term.

7. Network Architectures and Parameters

In this section we give all the details about the architectures and the parameters used in various experiments shown in the main paper.

7.1. Capturing Clustered Modes: Toy Example

As mentioned in the paper, in this experiment we used 'generator identification' based diversity enforcing term

for the MAD-GAN (our approach). The architecture and dataset details are given below.

Architecture Details: The architecture used is similar to the one used by [4] for their synthetic experiment. The generator network has 2 ReLU hidden layers with 128 neurons and batch normalization. In the case of MAD-GAN, we used three generators with the first layer being shared among them. It generates 2D samples in this case. Similar to [4], the input to each of the generator is a 3D normal noise $N(0, I_3)$. The discriminator consists of one fully connected layer of 128 ReLU neurons, mapping the 2D point to the output layer.

In the case of GAN and GAN-BN (GAN with batch normalization), the discriminator output is a scalar value between 0 to 1 obtained using the sigmoid. The loss thus is based on the binary cross entropy. While for our proposed MAD-GAN, dimension of the output layer is $k + 1 = 4$, which is normalized using a softmax layer. Thus, the loss in this case is the multi-label cross entropy. In the case of GAN, batch normalization is removed from the Generator. For the training, we used Adam optimizer with batch size of 128 and learning rate of for $2e-4$ for both generator and discriminator.

Dataset Generation We generated synthetic data with 3 clustered modes. To do so, we used GMM with 18 Gaussians and select their means such that it forms three clusters each with 6 Gaussians. The standard deviation used is 0.1, similar to [4]. The means of the Gaussians were such that they formed a regular hexagon. The centers of each of these three hexagons lied on the vertices of an equilateral triangle.

7.2. Image to Image Translation

Architecture details: The network architecture is adapted from [11] and the experiments were conducted with the U-Net architecture and patch based discriminator.

In more detail, let Ck denote a Convolution-BatchNorm-ReLU layer with k filters and CDk represent a Convolution-BatchNorm-Dropout-ReLU layer with a dropout rate of 50%. All Convolutions are 4x4 spatial filters with a stride of 2. Convolutions in the encoder, and in the discriminator, downsample by a factor of 2, whereas in the decoder they upsample by a factor of 2.

Generator Architectures We used the U-Net generator based architecture from [11] as follows:

- U-Net Encoder: C64-C128-C256-C512-C512-C512-C512-C512
- U-Net Decoder: CD512-CD1024-CD1024-C1024-C1024-C512-C256-C128. Note that, in case of MAD-GAN, the last layer does not share parameters with other Generators.

After the last layer in the decoder, a convolution is applied to map to the number of output channels to 3, followed by a Tanh function. BatchNorm is not applied to the first C64 layer in the encoder. All ReLUs in the encoder are leaky, with a slope of 0.2, while ReLUs in the decoder are not leaky. The U-Net architecture has skip-connections between each layer i in the encoder and layer $n - i$ in the decoder, where n is the total number of layers. The skip connections concatenate activations from layer i to layer $n - i$. This changes the number of channels in the decoder.

Discriminator Architectures The patch based 70x70 discriminator architecture was used in this case : C64-C128-C256-C512.

Diversity terms

- Similarity based competing objective: After the last layer, a convolution is applied to map to a 1 dimensional output followed by a Sigmoid function. For the unsupervised feature representation $\phi(.)$, the feature activations from the penultimate layer C256 of the discriminator was used as the feature activations for the computation of the cosine similarity.
- Generator identification objective: After the last layer, a convolution is applied to map the output layer to the dimension of $k+1$ (number of generators + 1) followed by the softmax layer for the normalization.

For the training, we used Adam Optimizer with learning rate of $2e-4$ (for both generator and discriminator), $\lambda_{L1} = 10$ (hyperparameter corresponding to the L_1 regularizer), $\lambda = 1e-3$ (corresponding to the similarity based competing objective), and batch size of 1.

Dataset Preparation:

- Edges To Handbags: We used 137k Amazon Handbag images from [27]. The random split into train and test was kept the same as done by [27].
- Night To Day: We used 17823 training images extracted from 91 webcams. We thank Jun-Yan Zhu for providing the dataset.

7.3. Face generation with InfoGAN as the generator

Architecture details: The architecture uses the one employed in the InfoGAN paper [5]. As in the InfoGAN paper, the number of sets of categorical variables used was 10 (each of dimension 10). These 100 variables formed the set of salient variables which took part in the formulation for

maximizing mutual information. Thus, remaining 148 variables were sampled from random noise. The discriminator D and the recognition network Q are described in the Table 1 while generator is described in the Table 2.

Discriminator D / recognition network Q
Input 32x32 Color Image
4x4 conv. 64 leakyRELU. stride 2
4x4 conv. 128 leakyRELU. stride 2. batchnorm
4x4 conv. 256 leakyRELU. stride 2. batchnorm
FC. output layer for D FC.128-batchnorm-leakyRELU-FC.output for Q

Table 1. The last layer output is 1 dimensional (normalizer is sigmoid) for the similarity based competing objective and $k+1$ dimensional (normalizer is softmax) for the generator identification based objective.

Generator G
Input $\in \mathbb{R}^{228}$
FC. 2x2x448 RELU.batchnorm.shared
4x4 upconv. 256 RELU.stride 2.batchnorm.shared
4x4 upconv. 128 RELU.stride 2.shared
4x4 upconv. 64 RELU.stride 2.shared
4x4 upconv. 3 Tanh.stride 2.

Table 2. All the layers except the last one are shared among all the generators.

For the training, we used Adam Optimizer with learning rate of $2e-4$ (both generator and discriminator) and batch size of 64.

Dataset preparation: We used CelebA dataset [16] consisting of $202,599$ celebrity face images. We used all of them for the unsupervised training.

7.4. Face/Image generations with DC-GAN

Architecture details: Our architecture uses the one employed for the face generation with DC-GAN [21]. Concretely, the discriminator architecture is described in Table 3 and the generator architecture in Table 4.

Discriminator D
Input 64x64 Color Image
4x4 conv. 64 leakyRELU. stride 2. batchnorm
4x4 conv. 128 leakyRELU. stride 2. batchnorm
4x4 conv. 256 leakyRELU. stride 2. batchnorm
4x4 conv. 512 leakyRELU. stride 2. batchnorm
4x4 conv. output leakyRELU. stride 1

Table 3. The last layer output is 1 dimensional (normalizer is sigmoid) for the similarity based competing objective and $k+1$ dimensional (normalizer is softmax) for the generator identification based objective.

Generator G
Input $\in \mathbb{R}^{100}$
4x4 upconv. 512 RELU.batchnorm.shared
4x4 upconv. 256 RELU. stride 2.batchnorm.shared
4x4 upconv. 128 RELU. stride 2.batchnorm.shared
4x4 upconv. 64 RELU. stride 2.batchnorm.shared
4x4 upconv. 3 tanh. stride 2

Table 4. All the layers except the last one are shared among all the three generators.

Diversity terms For the 'competing Objective' based on the similarity function, the feature activations of the second last layer was used as $\phi(.)$ for the cosine similarity based computation.

For the training, we used Adam Optimizer with the learning rate of $2e-4$ (both generator and discriminator), $\lambda = 1e-4$ (competing objective), and batch size of 64.

Dataset preparation: We used CelebA dataset as mentioned in Section 7.3 for Face generation based experiments. For Image generation all the images (14,197,122) from the Imagenet-1k dataset [6] were used to train the DCGAN with 3 Generators alongside the 'generator identification' based objective. The images from both CelebA and Imagenet-1k were resized into 64x64.

7.5. Unsupervised Representation Learning

Architecture details: Our architecture uses the one proposed in DCGAN [21]. Similar to the DCGAN experiment on SVHN dataset (32x32x3) [19], we removed the penultimate layer of generator (second last row in Table 4) and first layer of discriminator (First convolution layer in Table 3).

Technique	2G	3G	4G
SCO	20.2%	19.6%	18.3%
GIO	20.5%	18.2%	17.5%

Table 5. Using similarity based competing objective (SCO) and generator identification based objective (GIO), the results of MAD-GAN on SVHN with different number of generators are shown.

Classification task: We trained our model on the available SVHN dataset [19]. For feature extraction using discriminator, we followed the same method as mentioned in the DCGAN paper [21]. The features were then used for training a regularized linear L2-SVM. The ablation study is presented in Table 5

Dataset preparation: We used SVHN dataset [19] consisting of $73,257$ digits for the training, $26,032$ digits for the testing, and $53,1131$ extra training samples. As done in DCGAN [21], we used 1000 uniformly class distributed random samples for training, $10,000$ samples from the non-extra set for validation and 1000 samples for testing.

For the training, we used Adam Optimizer with learning rate of $2e-4$ (both generator and discriminator), $\lambda = 1e-4$ (competing objective), and batch size of 64.

References

[1] M. Abadi and D. Andersen. Learning to protect communications with adversarial neural cryptography. *arXiv preprint arXiv:1610.06918*, 2016. 2

[2] M. Arjovsky and L. Bottou. Towards principled methods for training generative adversarial networks. Technical report, 2017. 1, 2, 3

[3] C. Bishop. Pattern recognition and machine learning (information science and statistics), 1st edn. 2006. corr. 2nd printing edn. *Springer, New York*, 2007. 7

[4] T. Che, Y. Li, A. Jacob, Y. Bengio, and W. Li. Mode regularized generative adversarial networks. Technical report, 2017. ICLR. 1, 2, 3, 10

[5] X. Chen, Y. Duan, R. Houthooft, J. Schulman, I. Sutskever, and P. Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. *arXiv preprint arXiv:1606.03657*, 2016. 1, 2, 3, 6, 7, 10

[6] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*, 2009. 6, 11

[7] I. Durugkar, I. Gemp, and S. Mahadevan. Generative multi-adversarial networks. *arXiv preprint arXiv:1611.01673*, 2016. 2

[8] A. Ghosh, V. Kulharia, A. Mukerjee, V. Namboodiri, and M. Bansal. Contextual rnn-gans for abstract reasoning diagram generation. In *AAAI*, 2017. 1

[9] A. Ghosh, V. Kulharia, and V. Namboodiri. Message passing multi-agent gans. *arXiv preprint arXiv:1612.01294*, 2016. 2

[10] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, 2014. 1, 2, 6

[11] P. Isola, J.-Y. Zhu, T. Zhou, and A. Efros. Image-to-image translation with conditional adversarial networks. *arxiv*, 2016. 2, 6, 7, 10

[12] T. Joachims, T. Finley, and C. Yu. Cutting-plane training of structural SVMs. *Machine Learning*, 2009. 4

[13] D. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 1

[14] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi. Photo-realistic single image super-resolution using a generative adversarial network. *arXiv preprint arXiv:1609.04802*, 2016. 1

[15] M. Liu and O. Tuzel. Coupled generative adversarial networks. *arXiv preprint arXiv:1606.07536*, 2016. 2

[16] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *(ICCV)*, 2015. 6, 11

[17] L. Metz, B. Poole, D. Pfau, and J. Sohl-Dickstein. Unrolled generative adversarial networks. In *ICLR*, 2017. 1, 2, 3

[18] M. Mirza and S. Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014. 1, 7

[19] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Y. N. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011. 2, 8, 11, 12

[20] D. Pathak, P. Krahenbuhl, J. Donahue, T. Darrell, and A. Efros. Context encoders: Feature learning by inpainting. *arXiv preprint arXiv:1604.07379*, 2016. 1

[21] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015. 2, 6, 7, 8, 11, 12

[22] S. Reed, Z. Akata, X. Yan, L. Logeswaran, B. Schiele, and H. Lee. Generative adversarial text to image synthesis. *arXiv preprint arXiv:1605.05396*, 2016. 1

[23] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training gans. *arXiv preprint arXiv:1606.03498*, 2016. 1, 2, 3, 4

[24] I. Tsochantaridis, T. Hofmann, T. Joachims, and Y. Altun. Support vector machine learning for interdependent and structured output spaces. In *ICML*, 2004. 4

[25] C. Vondrick, H. Pirsiavash, and A. Torralba. Generating videos with scene dynamics. *arXiv preprint arXiv:1609.02612*, 2016. 1

[26] J. Wu, C. Zhang, T. Xue, W. Freeman, and J. Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. *arXiv preprint arXiv:1610.07584*, 2016. 1

[27] J. Zhu, P. Krähenbühl, E. Shechtman, and A. Efros. Generative visual manipulation on the natural image manifold. In *ECCV*, 2016. 1, 10